arXiv:1608.08711v1 [cs.CV] 31 Aug 2016

# Engagement Detection in Meetings *

Maria Frank[1], Ghassem Tofighi[2], Haisong Gu[3], and Renate Fruchter[4]

[1] Civil and Environmental Eng., Stanford University,
Stanford, California, U.S.A.
mfrank0@stanford.edu
[2] Electrical and Computer Eng., Ryerson University,
Toronto, Ontario, Canada
gtofighi@ryerson.ca
[3] Machine Learning Devision, Konica Minolta Laboratory U.S.A.,Inc
San Mateo, California, U.S.A.
haisong.gu@hl.konicaminolta.us
[4] Project Based Learning Laboratory (PBL Lab),
Civil and Environmental Eng., Stanford University,
Stanford, California, U.S.A.
fruchter@stanford.edu

**Abstract**

Group meetings are frequent business events aimed to develop and conduct project work, such as Big Room design and construction project meetings. To be effective in these meetings, participants need to have an engaged mental state. The mental state of participants however, is hidden from other participants, and thereby difficult to evaluate. Mental state is understood as an inner process of thinking and feeling, that is formed of a conglomerate of mental representations and propositional attitudes. There is a need to create transparency of these hidden states to understand, evaluate and influence them. Facilitators need to evaluate the meeting situation and adjust for higher engagement and productivity. This paper presents a framework that defines a spectrum of engagement states and an array of classifiers aimed to detect the engagement state of participants in real time. The Engagement Framework integrates multi-modal information from 2D and 3D imaging and sound. Engagement is detected and evaluated at participants and aggregated at group level. We use empirical data collected at the lab of Konica Minolta, Inc. to test initial applications of this framework. The paper presents examples of the tested engagement classifiers, which are based on research in psychology, communication, and human computer interaction. Their accuracy is illustrated in dyadic interaction for engagement detection. In closing we discuss the potential extension to complex group collaboration settings and future feedback implementations.

**Keywords**: *Collaboration*, *Engagement*, *Feedback*, *Meeting Management*

## 1 Introduction

Project meetings, such as Big Room design and construction project sessions (Khanzode & Lamb, 2012 [1]), are communicative events where participants discuss, negotiate, present, and create material jointly. Participants express varying mental states in these group meetings due to individual factors and predispositions, the specific meeting agenda, and general team dynamics. To improve meeting quality and productivity it is important to understand the

*Empirical data collected in Machine Learning devision at Konica Minolta Laboratory U.S.A.,Inc. The algorithms and methods provided in this research are submitted to United States Patent and Trademark Office (USPTO).

individual user states and the aggregated group state and through this the group dynamics. One characteristic associated with meeting quality, productivity, and effectiveness is engagement.

This paper focuses on a multi-modal approach to detect engagement. Engagement is defined as: the value that a participant in an interaction attributes to the goal of being together with the other participant(s) and continuing interaction (Salam, Chetouani 2015, p. 3[2]) the process by which two (or more) participants establish, maintain, and end their perceived connection. It is directly related to attention (Salam, Chetouani 2015, p. 3[2]). It can furthermore be defined as a meaningful involvement (Fletcher 2005[3]), that is enabled through vigor, dedication, and absorption (Schaufeli 2006[4]). For the present study we define engagement as an attentive state of listening, observing, and giving feedback, leading into protagonistic action in group interaction. Individual and group engagement level influences productivity of group interaction. Hence, it is of interest to create transparency about engagement states of participants and group.

An important aspect is the intention to become protagonist and the identification of the threshold between active and passive participation. This moment is crucial because it indicates a significant change in mental state of an individual. Intention is understood as instructions that people give to themselves to behave in certain ways (Triandis 1980[5]), and one of the most important predictor of a persons behavior (Sheeran 2002[6]). It is an internal commitment to perform an action while in a certain mental state (Levesque et al. 1990[7]) and it predicts actual behavior. This paper defines intention as a mental state characterized by an internal commitment to perform a specific action in the immediate future. It is a specific form of engagement that is indicating future active meeting participation.

We build on an early study and preliminary prototype, called eRing (engagement Ring), at the PBL Lab at Stanford that focused on detecting and building awareness of learners degrees of engagement during globally distributed project team meetings (Ma & Fruchter, 2015[8]). Collocation offers a wealth of tacit cues about learners and team members. Body position and gestures, as well as physical movement provide indicators of participants intentions (Pease & Pease, 2004[9]). These are cues that enable us to intuitively interpret and evaluate the state of learners and team members.

Integrated Delivery Process (IPD) is one of the disruptive forces that the design and construction industry is experiencing. This process leverages Big Room weekly or bi-weekly collocated project stakeholder meetings. During Big Room sessions a lot of time is spent talking in large groups about issues, displaying side by side on SMARTBoards 3D BIM models of the future facility and the Navisworks clashes that need to be resolved (Khanzode & Lamb, 2012[1]). Big Room IPD increases the number of iterations and time-to-market (Fruchter & Ivanov, 2011[10]). In order to achieve this increased productivity participants need to engage and get feedback on their engagement state.

This paper presents a more comprehensive Engagement Framework that builds on and expands the eRing engagement concept and provides guidelines and instructions for the implementation of a multi-modal sensor-based work environment to detect engagement of participants in a collocated setting. In the preliminary eRing prototype the initial focus was on body motion and posture. The proposed Engagement Framework presented in this paper expands this approach to include other information streams such as facial expression, voice, and other biometric information. We demonstrate the accuracy of the Engagement Framework with a simple interaction scenario, and provide a direction for future implementation.

## 2 Theoretical Points of Departure

The analysis of engagement has a long history in education and work. Work engagement (Schaufeli 2006[4]) influences project outcome. Qualitative observation-based evaluation of an individuals state has shed light on the importance of engagement for productivity. The Human-Computer-Interaction community started focusing on the automation of engagement detection for learning and games. To combine both methods of engagement detection, qualitative and automated, enables us to provide new insights into automated work engagement detection in the work environment.

A study by Frank and Fruchter (2015[11]) uses cognitive flexibility as an internal evaluation of engagement during meetings. However, this approach as well as self-reports require active involvement of meeting participants. This paper focuses on external expressions of engagement that allow a non-intrusive evaluation of the participant without disrupting the activity. Research in this field focuses on the human body and its features to generate indicators of mental states. Non-verbal behavior and body language expresses emotions (Schindler et al. 2008, Mead 2011[12]); it conveys information about the speaker's internal state, and attitude toward the addressees (Gutwin & Penner 2002[13]). Both body posture and gesture (Ekman 1999[14], Ekman and Friesen 1972[15]) have been identified as important social signals that indicate mental states and thereby also engagement. We use the preexisting qualitative research to develop an Engagement Framework offering quantitative measureable indicators.

Various approaches have been investigated to use body language to improve interactions with digital agents such as assistant robots and displays (Vaufrydaz 2015[16], Schwarz 2014[17]). These studies use varying aspects of non-verbal language, such as upper body pose (Mead 2011[12], Schwarz 2014[17]), proxemics (Vaufreydaz 2015[16]), facial expression and head posture (Vaufreydaz 2015[16]), or the absence of movement (Witchel 2014 [18]) and sometimes the combination of a small set of features from facial expression and posture in multi-modal approaches.

Weight and weight distribution on a seat has been used to measure engagement (Mota & Picard 2003[19], DeMello et al. 2007[20], Balaban 2004 [21]). Weight is an indicator of body posture that can be measured without affecting the participant.

The aforementioned eRing prototype (Ma & Fruchter 2015 [8] focuses on globally distributed learning settings which strip away the physical cues. The hypothesis of the eRing study was that these cues can be reintroduced in virtual interactions, allowing learners to self-regulate and build an awareness of the overall state of the team. The PBL Lab researchers developed a cloud service and application called eRing (engagement Ring) to detect and provide real-time feedback of learners degree of engagement during a virtual interaction. eRing collects body motion data using Microsoft Kinect sensor, analyzes and interprets a small set of the body motions and body positions including head, shoulder, and hand joints. eRing provides real-time feedback of the degree of engagement based on three states disengaged, neutral listen, and engaged. The body motion analysis is performed for two units of analysis individual learner and team. eRing runs on a Microsoft Azure cloud platform. eRing was used in the architecture, engineering, construction (AEC) Global Teamwork course testbed in 2014 and 2015.

However, the afore mentioned studies discuss limited sets of potential engagement indicators and do not make full usage of the amount of qualitative research conducted on non-verbal body language and its indication of mental states. Hence, they all provide limited insight and accuracy in evaluating the intention to engage with an agent. They do not discuss the evaluation of engagement between humans in normal interaction or insights into potentially different levels of engagement and thereby provide only a binary classification. This paper discusses a more

comprehensive analysis of engagement.

## 3 Proposed Engagement Classification Framework

Building on previous studies, we developed an Engagement Framework in the form of an engagement scale for interaction. Engagement is the level of participation in a meeting interaction. Schwarz et al. (2014)[17] define two states of disengaged and intention to interact. The eRing prototype defines three states, disengaged, neutral listen, and engaged. The proposed Engagement Framework consists of six states to understand engagement in more detail. Figure 1 presents the Engagement Framework consisting of six states. The first state is disengagement; similar to Schwarz et al. (2014)[17] and Ma and Fruchter (2015)[8] disengagement is understood as a state of no participation, distraction and no attention to the meeting. The second state is relaxed engagement, understood as attention to the meeting, listening, observing, but no participation. This is most similar to the neutral listen state of eRing. The third state is involved engagement it signifies attention and non-verbal feedback like nodding, small verbal confirmations, or specific facial expressions such as an open mouth. This state is the first that can be classified as engaged in the eRing prototype, but counts as disengaged in the Schwarz et al. discussion. The present framework differentiates between relaxed and involved engagement because they signify differing mental states and different focus and involvement to the meeting activity and thereby potential for productivity. The fourth state is the intention to act which is similar to the Schwarz model. It is the preparation for active participation in protagonist role indicated through and increase in activity. It shows an involved participation but not yet a contribution. The fifth state is action and the process of speaking and/ or interacting with participants or content on displays. It is a calm form of protagonistic activity. The sixth and final state is involved action and a highly engaged and involved interaction with intense gesture and voice. This state is an active form of excitement and arousal in a meeting that fosters productivity. The six states help understand different team dynamics and user states that affect the meeting interaction, productivity and effectiveness. Using multi-modal information our framework identifies specific combination of engagement classifiers associated with the six engagement states.

## 4 Engagement Classifiers

Literature identified specific observable patterns that indicate mental states such as engagement. In order to categorize the engagement of an individual into one of the six engagement states of the Engagement Framework we need these specific observable patterns. They elucidate which state a participant is in at any given type.

The Engagement Framework draws on a variety of these identified observable patterns understood as indicators for engagement (Sanghvi et al. 2011[22], Vaufrydaz et al. 2015[16], Mead et al. 2011[12], Bassetti 2015[23], Changing minds 2015, Schwarz 2014[17], Scherer 2011[24]). As humans we draw on a multitude of indicates to evaluate conversation partners engagement. For this reason, the Engagement Framework seeks to implement many identified indicators of engagement to make a more accurate prediction of the engagement state. The indicators are multi-modal, based on body posture and motion, voice, and facial expression to create a more comprehensive perspective. The multi-modal indicators can be captured through ubiquitous sensors such as 3D cameras like Kinect, 2D imaging through RGB cameras, and microphones. The input of each sensor and each mode (3D, 2D, sound) is treated as a module. The module

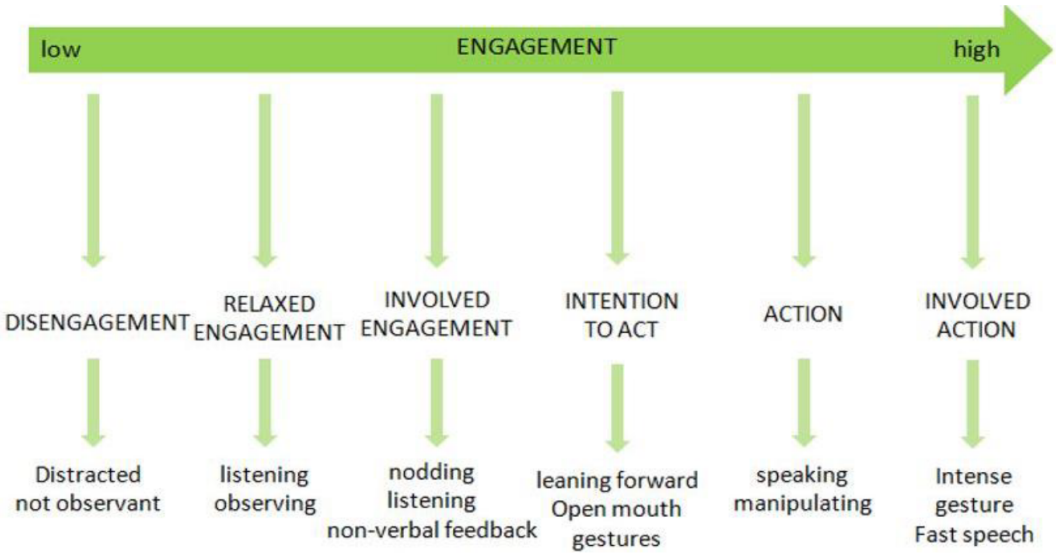


Figure 1: Engagement Framework and sequence of states

conceptualization allows for future expansion of sensors and indicators to be integrated into the classification of the Engagement Framework.

The indicators are abstracted into classifiers. Classifiers are specific components of body posture or motion, facial expression, voice, and other multi-modal information that can characterize engagement. The classifiers are based on heuristics. To date, we identified 25 indicators, such as volume of speech, smiling, and shaking the head, producing a range of classifiers. Participants in meetings are normally seated behind a table. Hence, the indicators are focused on the upper body with a differentiation between head, face, arms, body, and voice. Table 1 shows an example of indicators and their classifiers for the Body. It names the indicator, its definition, the module, and the classifier in each of the six engagement states. Classifiers are understood as being exhibited or not exhibited during an observation period and are expressed as a binary value. For example, the body lean angle is an indicator of engagement that contains six classifiers. The first classifier is a backwards lean angle that can indicate disengagement. The second classifier is a sideways lean angle that indicates relaxed engagement. The third classifier is no lean angle or sitting straight which shows an involved engagement. The fourth classifier is the process of leaning forward indicating the intention to act. The fifth classifier is a slight forward lean indication action. And the sixth classifier is a rapid change in lean angles over the observation time indicating involved action.

# 5 Implementation

Schwarz et al. (2014[17]) used a linear regression to calculate weight factors to evaluate the relative importance of five binary classifiers and defined a threshold for the sum of the classifiers to evaluate intention to act. This method is not a valid statistical approach. We propose the implementation of the classifiers as a feature vector containing the binary values. Depending on the available sensors the feature vector consists of features from different modules. The feature

Table 1: Example of classifier selection and definition for the Body

| indicator | definition | module | dis-engagement | relaced engagement | involved engagement | intention to act | action | involved action |
|---|---|---|---|---|---|---|---|---|
| **BODY** | | | | | | | | |
| lean angle | leaning in or back for attention | 3D | back against the chair, potentially to one side | towards one side | straight up, slight lean to the side | leaning forward | forward, or straight up | straight up, or multiple changes |
| slouch factor | tiredness | 3D | heavy slouch | medium slouch | slight slouch | upright | upright | upright |
| weight distribution | related to lean angle and weight on seat | weigth | weight on back of seat | on back or middle of seat | in middle | front of seat | front to middle of seat | middle, changing |

vector is fed to a Support Vector Machine with a linear kernel as a machine learning algorithm to evaluate the important classifiers and cluster the emerging pattern in the Engagement Framework states. Our approach is based on LIBSVM. LIBSVM implements the Sequential minimal optimization (SMO) algorithm for kernelized SVMs, supporting classification and regression.

Figure 2 illustrates the first stage of an example implementation. Each meeting includes various team members. For the implementation system these team members are unclassified (Figure 2a). Each individual meeting participant is represented through one feature abstraction, which is a skeleton 3D sensor abstraction in this example (Figure 2b). Based on the skeleton the related set of binary classifiers is instantiated and creates the feature vector for each team member (Figure 2c). The SVM then classifies the individual engagement state from one to six (disengagement to involved action) accordingly (Figure 2d).

Figure 2: Implementation of Engagement Framework and evaluation of engagement state

Figure 3 illustrated the final team engagement classification. The Engagement Framework Implementation creates multiple feature vectors to evaluate every participant in a meeting situation and classify their indicators into one of the six engagement states 3a). The classification result for each team member is then aggregated to show the engagement state distribution of the team and calculate the average engagement state as team state (Figure 3b).

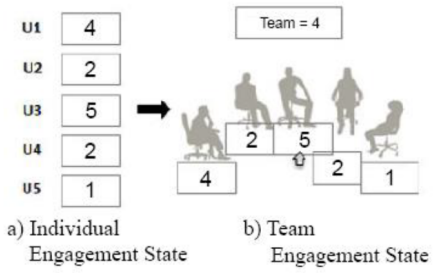


Figure 3: Classification of individual and team engagement state

# 6 Testing

## 6.1 Scenario

In an initial experiment a dyad is observed to train the algorithm and test prediction accuracy. The focus of this experiment is, similar to eRing (Ma & Fruchter 2015[8]) on three out of the six states: disengagement, intention to act, action (Figure 4). We implemented the body motion and posture module based on 3D joint information for this experiment. Due to the complexity of the Engagement Framework, and the implementation through multi-modal modules and classifiers we had to simplify the first experiment as a starting point. This simplification allows for an initial engagement classification and testing of the implementation of the classifiers.

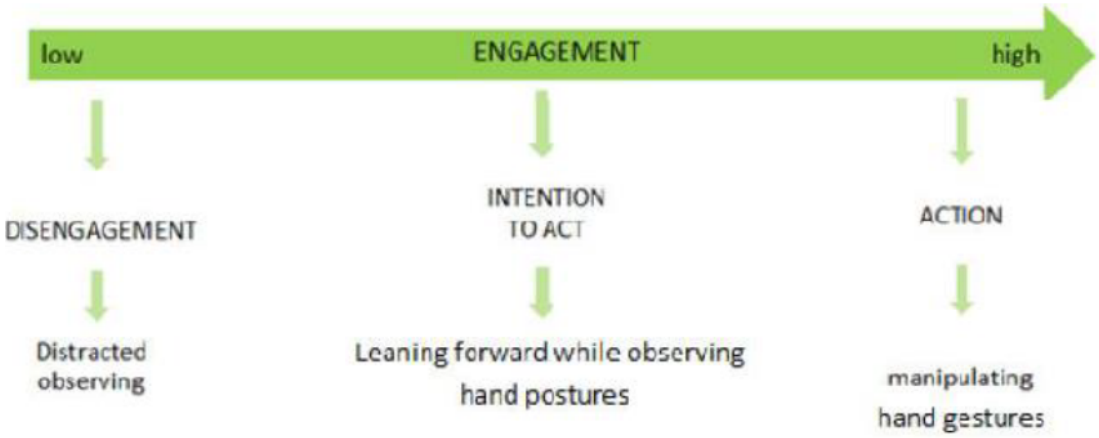


Figure 4: States implemented in initial experiment

## 6.2 Data Collection

Schwarz (2014[17]) introduced a simple task hand-off game to evaluate intention to interact. We implemented a similar task hand-off game to observe and detect the three engagement states of disengagement, intention to act, and action and create a clear ground truth. The dyads task is

to select colored blocks on a display by using a hand to navigate a cursor. Colors of the blocks change for different participants. When participant 1 is playing the blocks are green. The blocks turn red when participant 2 is playing. The task starts with a 10 second countdown, participant 1 then has 10 seconds toselect as many blocks as possible. There is a 10 second switch period for participant 2 to get ready and have the intention to act. Participants are penalized for not disengaging after their 10 second segment. The game runs for 20 total switches. Each participant group plays the game 5 times. Participants were positioned in about 2.5 m distance from a screen. A depth-sensing camera Kinect (Xbox one) is installed underneath the screen to capture body motion.

Data was collected from six members of the research team at Konica Minolta: 3 female, 3 male, all are professionals in IT and have a moderate to high experience with motion detection. The participants are randomly assembled in pairs to play the game.

### 6.3 Implemented Classifiers

The 3D sensor module using the ASUS Xtion, used the data of the upper body joints (Figure 5). Participants are normally seated behind a table in meetings. Therefore, only the body joints for head, left and right shoulders, left and right elbows, left and right hand, torso and left and right hips are relevant.

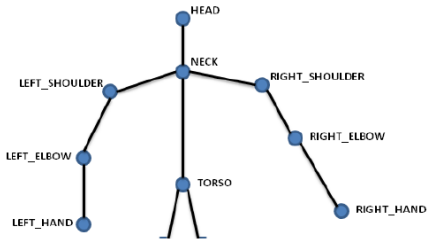


Figure 5: States implemented in initial experiment

Table 2 shows the list of binary classifiers implemented in the initial experiment. A total of 5 indicators and their 16 classifiers is used to create the feature vector. The classifiers are evaluated on a frame by frame basis. This means that for each frame of the task recording a feature vector of the binary classifiers is created and the engagement state classified.

Table 2: Classifiers of initial experiment

| Indicator | Classifiers |
|---|---|
| hand vertical position | below hip, below torso, belor shoulder, below head |
| hand speed | stopped, slow, fast, very fast |
| body direction | facing sensor |
| leaning | lean back, no lean, lean forward |
| special posture | hands folded, hands on head, arms crossed |

We captured and labeled 2,321 frames from 6 different subjects for Intention to Act and

Disengagement. Five hundred frames (about 20% of the frames) are used for training and the remaining frames are used for testing our classifiers. Each frame is classified as Disengagement or Intention to Act. The Action state occurs when a frame label is Intention to Act and the speed of one hand is not zero. The initial implementation reached a classification accuracy of 83.36%.

The processing time for each frame is less than 10ms which indicates real-time usability of our algorithm since we receive 30 frames per second and should label them to one of the states of engagement.

# 7 Discussion

Automated engagement detection can assist to increase our understanding of group dynamics, and participation habits. The presented Engagement Framework was implemented using research based observable engagement indicators and classifiers show a high accuracy in predicting participation behavior. This provides the basis for future comprehensive team engagement analyses. Our proposed approach to engagement classification promises a more comprehensive observation of individual engagement and team engagement for various applications.

Throughout the meeting the participation level of participant changes. Calculating the engagement state of individual and group in real time allows for feedback of group activity and enables managerial and environmental adjustment to create a more engaging and thereby productive work environment.

## 7.1 Engagement Feedback For Meeting Participant and Infrastructure Management

The individual and team level engagement state can help to inform the meeting in process on various levels, and we elaborate on some of the application opportunities.

A dynamic meeting feedback can trigger a signal at individual or team level; either targeting the disengaged individual to encourage reengagement or the group to change the meeting focus or activity. Groups would receive a joint message such as The group is disengaged or a motion signal is triggered like eRing (Ma & Fruchter 2015[8]), which means that the SVM classified more than 40% of the individuals as disengaged. The public feedback is used as an encouragement to change the meeting activity or take a break to refocus. Detailed insight is helpful to evaluate meeting effectiveness and efficiency.

Records of engagement levels will provide insight to bigger trends in meeting dynamic, such as overall engagement at specific meeting times, participation, and dominance in meetings. These insights allow for strategic adjustments of meeting times, necessary participants, necessary changes in interaction protocols, and others aspects related to meeting productivity.

## 7.2 Operation Intent Feedback

The engagement level can be used to identify the user intent to act with various responsive objects in a room by identifying directed engagement levels. It will give additional information to responsive objects as to which user wants to interact with which object in a room, based on directional body posture, voice, or motion. It will use an engagement table to track each participants potential engagement level with each object.

### 7.3 Limitations and Future Research

The presented paper is an initial step in developing a complex responsive engagement system. The accuracy of the algorithm can be increased. In this initial stage of the project a limited number of classifiers have been implemented. Currently the study focuses on body posture and motion. The classifiers are designed as modules that allow for expansion through different sensors, integrating for example voice data, and facial expression, up to visions of the smart office future including biometric information. The expansion of the classifier list will allow for a more complex understanding of engagement and user state. Furthermore, on-going efforts focus on the implementation of a machine learning algorithm that will allow for a faster and more accurate classification of the six engagement states. The collection of more training data will further assist is making the Framework implementation more robust.

The test game presented in the paper is an abstraction and simplification of real interactions and hence future experimental scenarios will consider more complex future steps. We will present the teams with increasingly complex interaction tasks up to the implementation of the algorithm in realistic meetings.

The final goal for the implementation of the Engagement Framework is to continuously classify the engagement of multiple meeting participants that are within the operating range of the sensors for all six engagement states. Throughout a meeting participants listen, observe, talk, or interact with a display. The behaviors are evaluated through multi-modal modules of engagement classifiers to categorize engagement into six stages. Such a system will assist to detect and evaluate engagement states, and foster engagement in scenarios ranging from small meeting groups to Big Room project stakeholder meetings.

## References


[1] Atul Khanzode and Eric Lamb. Transcending the bim hype, 2012.

[2] Hanan Salam and Mohamed Chetouani. A multi-level context-based modeling of engagement in human-robot interaction. In *Automatic Face and Gesture Recognition (FG), 2015 11th IEEE International Conference and Workshops on*, volume 3, pages 1–6. IEEE, 2015.

[3] Adam Fletcher. Meaningful student involvement. *Retrieved on December*, 1:2012, 2012.

[4] Wilmar B Schaufeli, Arnold B Bakker, and Marisa Salanova. The measurement of work engagement with a short questionnaire a cross-national study. *Educational and psychological measurement*, 66(4):701–716, 2006.

[5] Harry C Triandis. Values, attitudes, and interpersonal behavior. In *Nebraska symposium on motivation*. University of Nebraska Press, 1979.

[6] Paschal Sheeran. Intentionbehavior relations: A conceptual and empirical review. *European review of social psychology*, 12(1):1–36, 2002.

[7] PR Cohen and HJ Levesque. Intention is choice with commitment artificial intelligence, 42 213261. *Cohen, PR and Levesque, HJ*, 1995.

[8] Ma, J and Fruchter, Renate. eRing: Body Motion Engagement Detection and Feedback in Global Teams. SAVI Symposium on New ways to teach and learn for student engagement, Stanford University, 2015. April 2015.

[9] Barbara Pease and Allan Pease. *The definitive book of body language*. Bantam, 2008.

[10] Renate Fruchter and Plamen Ventsislavov Ivanov. Agile ipd production plans as an engine of process change. In *Proc. ASCE International Workshop on Computing in Civil Engineering*, pages 776–784, 2011.

[11] Maria Frank, Renate Fruchter, and Marianne Leinikka. High engagement decreases fatigue effect in global learners. In *SAVI Symposium on New ways to teach and learn for student engagement*. Stanford University, 2015.

[12] Ross Mead, Amin Atrash, and Maja J Matarić. Proxemic feature recognition for interactive robots: automating metrics from the social sciences. In *International conference on social robotics*, pages 52–61. Springer, 2011.

[13] Carl Gutwin and Reagan Penner. Improving interpretation of remote gestures with telepointer traces. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work*, pages 49–57. ACM, 2002.

[14] Paul Ekman. Emotional and conversational nonverbal signals. In *Language, knowledge, and representation*, pages 39–50. Springer, 2004.

[15] Paul Ekman and Wallace V Friesen. Hand movements. *Journal of communication*, 22(4):353–374, 1972.

[16] Dominique Vaufreydaz, Wafa Johal, and Claudine Combe. Starting engagement detection towards a companion robot using multimodal features. *Robotics and Autonomous Systems*, 75:4–16, 2016.

[17] Julia Schwarz, Charles Claudius Marais, Tommer Leyvand, Scott E Hudson, and Jennifer Mankoff. Combining body pose, gaze, and gesture to determine intention to interact in vision-based interfaces. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 3443–3452. ACM, 2014.

[18] Harry J Witchel, Carina EI Westling, Julian Tee, Aoife Healy, Robert Needham, and Nachiappan Chockalingam. What does not happen: Quantifying embodied engagement using nimi and self-adaptors. *Participations*, 11(1):304–331, 2014.

[19] Selene Mota and Rosalind W Picard. Automated posture analysis for detecting learner's interest level. In *Computer Vision and Pattern Recognition Workshop, 2003. CVPRW'03. Conference on*, volume 5, pages 49–49. IEEE, 2003.

[20] Sidney DMello, Patrick Chipman, and Art Graesser. Posture as a predictor of learners affective engagement. In *Proceedings of the 29th annual cognitive science society*, volume 1, pages 905–910. Citeseer, 2007.

[21] Carey D Balaban, Joseph Cohn, Mark S Redfern, Jarad Prinkey, Roy Stripling, and Michael Hoffer. Postural control as a probe for cognitive state: Exploiting human information processing to enhance performance. *International Journal of Human-Computer Interaction*, 17(2):275–286, 2004.

[22] Jyotirmay Sanghvi, Ginevra Castellano, Iolanda Leite, André Pereira, Peter W McOwan, and Ana Paiva. Automatic analysis of affective postures and body motion to detect engagement with a game companion. In *2011 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 305–311. IEEE, 2011.

[23] Chiara Bassetti. Social gatherings: Integrating sociological theory and analysis into computer vision methods and techniques. *Workshop on Group and Growd Behavior Analysis and Understanding*, 2015.

[24] Stefan Scherer, Michael Glodek, Georg Layher, Martin Schels, Miriam Schmidt, Tobias Brosch, Stephan Tschechne, Friedhelm Schwenker, Heiko Neumann, and Günther Palm. A generic framework for the inference of user states in human computer interaction. *Journal on Multimodal User Interfaces*, 6(3-4):117–141, 2012.